A Semi-Autonomous Weapon Payload

G. Kogut[a], L. Drymon[a] ,H.R. Everett[a], E. Biagtan Pacis[a] ,H. Nguyen[b], B. Stratton[a]
J. Goree[c], B. Feldman[c]

[a]Space and Naval Warfare Systems Center, San Diego *
[b]Science Applications International Corporation
[c] Telerobotics Corporation

ABSTRACT

Weapon payloads are becoming increasingly important components of unmanned ground vehicles (UGVs). However weapon payloads are extremely difficult to teleoperate. This paper explores the issues involved with automating several aspects of the operations of a weapon payload. These operations include target detection, acquisition, and tracking. Various approaches to these issues are discussed, and the development and results from two different working prototype systems developed at Space and Naval Warfare Systems Center, San Diego (SSC San Diego) are presented. One approach employs a motion-based scheme for target identification, while the second employs an appearance based scheme. Target selection, arming and firing remain teleoperated in both systems.

Keywords: robotics, weapons, payload, tracking, aiming, sensors

1. INTRODUCTION

Small unmanned ground vehicles (UGVs) are becoming increasingly common on the battlefield. The first generation of deployed UGVs was limited to applications such as tunnel and cave exploration and explosive ordnance disposal (EOD). However, there are currently efforts underway to provide UGVs with weapon payloads. Examples include the Special Weapons Observation Reconnaissance Detection System (SWORDS), and the Gladiator Tactical Unmanned Ground Vehicle (TUGV). While larger, more powerful tactical UGVs are being developed under programs like Future Combat Systems (FCS), the first generation of tactical UGVs are generally inexpensive, lightweight UGVs armed with weapons that were designed for operation by a single human. For example, the SWORDS and Gladiator UGVs are designed to be fitted with the M240 Medium Machine Gun or similar weapons. These weapons are generally mounted on pan-tilt mounts, allowing the weapon a wide field of fire independent of the pose or movement of the UGV itself.

The weapons mounted on this first generation of tactical UGVs were designed to be controlled by teleoperation. Teleoperation is usually performed by a joystick or other two degree of freedom controller mounted in close proximity to the UGV's standard teleoperation controls. Feedback to the operator is produced by a range of sensors on the UGV, but primarily consists of visual and infrared

---

* Author Information: Email: greg.kogut@navy.mil, Phone: (619) 553-0707, Fax (619) 553-6188,
Web: http://www.spawar.navy.mil/robots/

cameras. The sensor data is displayed on small LCD screens. The screen and controls are generally packaged together in one man-portable, battery-powered unit

However, there are several limitations and drawbacks to this concept of weapon control on a UGV. These include the problems of 1) excessive burden on the operator, 2) control and feedback latencies, 3) limited situational awareness.

1.1 Burden on the Operator

Operating a teleoperated UGV is a difficult task. UGVs are complex machines, typically consisting of drive motors, a wide array of sensors, lights, actuators, and other components. Each of these components typically provides feedback to the operator (video, temperature, voltage level, etc), and requires an input from the user. Therefore the typical user control interface consists of at least two joysticks, or other two degree of freedom controllers, many buttons and switches, and at least one video screen. Many of these controllers serve multiple purposes, depending on the mode of operation. For example, a camera pan-tilt joystick can become a weapon aiming controller when the weapon system is activated. The end result is that such a controller is extremely complex and may require many hours of training before a user is competent and comfortable enough to use the system during a real mission.

These problems are exacerbated in a system with a weapon payload for several reasons, including: increased operator control unit (OCU) complexity, potential time constraints on weapon payload tasks, and the requirement to follow the mission rules of engagement while simultaneously controlling the weapon.

A weapon payload adds several more controls to an operator control unit. The typical weapon payload requires, minimally, a two degree of freedom controller, one or more arming buttons, and a fire button. Often there are many more controls for triggers safeties, pan and tilt rates and limits, etc. This additional complexity adds to the already crowded OCU of most UGVs.

UGV operators controlling a UGV on an EOD or tunnel exploration mission generally have no hard time constraint. They can work in a relaxed manner, at their own pace. By contrast, when the use of a weapon is necessary, there is often a very limited time to successfully complete a task before either the intended target escapes, or, if armed, successfully attacks the UGV. This time constraint greatly increases operator stress and the probability of operator error.

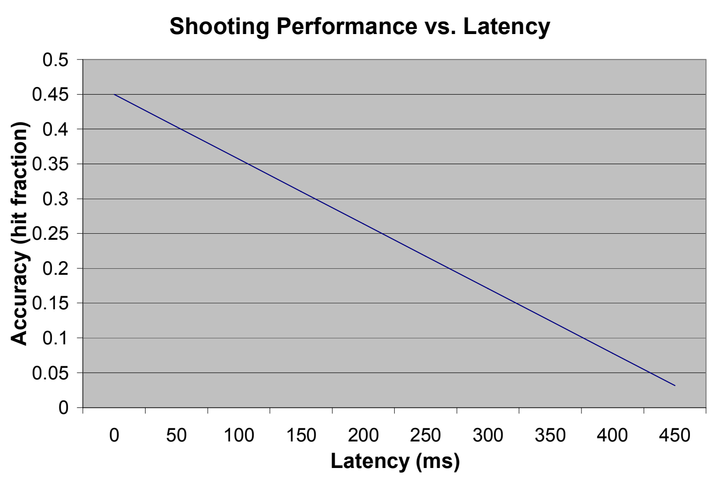

Figure 1. The effect end-to-end latency on manual shooting performance in the video game Unreal Tournament 2003[5] Latency approaching half a second essentially renders hitting a moving target impossible.

An operator of a weapon payload must also follow the rules of engagement before arming and firing a weapon. The operator is held responsible for the decision, just as if he were holding the weapon himself. These decisions must be made simultaneously with the control tasks listed above. This is an additional significant stress placed upon the operator.

SSC San Diego is working on several projects which aim to reduce the overall burden on the user of UGV's, including UGV's with weapon payloads

1.2 Feedback and Control Latency

There are two sources of latency end-to-end system in a teleoperated UGV. The first is sensor feedback latency. Teleoperated UGVs are primarily controlled by visual feedback, in which case the feedback latency is the latency in video transmission from a camera mounted on the UGV to the remote OCU's video display screen, including potential latencies in video digitizing, encoding, and decoding. The second form is control latency. This is the latency is the transmission time of a user's control response in addition to the time it takes the UGV to act on a received control response.

A UGV with a weapon payload will likely operate with a wireless communication link, and will operate outside the line-of-sight of the operator and OCU. This assumption is made because the primary motivation for developing weapon payloads is to remove humans from harm's way, which requires moving them from the line of sight of an enemy.

Therefore, the primary factor in system latency is the wireless communication link. Current non-line-of-sight wireless communication links generally do not support enough data throughput to carry full-size, real-time uncompressed video. Techniques to compress video (and then decompress at the OCU) can make non-line-of-sight video transmission possible, but these techniques typically increase latency. For example, required throughput can be reduced by decreasing video frame rates, but this approach introduces inherent latency equal (at least) to the time gap between frame transmissions plus the time to transmit a frame. Other approaches try to maintain reasonably high frame rates but use intra-frame and inter-frame compression to reduce bandwidth. Inter-frame compression, used in such compression standards as MPEG4 and H.263, necessarily adds latency to the process of compressing and decompressing video since decoding an individual frame requires data from subsequent video frames.

The combination of these latency sources can produce overall system latencies which range from several hundred milliseconds to tens of seconds. The effect of such latency in the teleoperation of remote vehicles is well documented. Elliott and Eagleson studied the effects of latency on teleoperation, and noted that, "…even latencies as a small as a few hundred milliseconds will prevent the operator from controlling a device in a natural way. Instead, the control of the remote system becomes difficult; it requires that the operator anticipate the effects of inaccuracies and unexpected events, which will not be known immediately because of the communications delay."[3] Boyle noted that delays in feedback caused operators to consistently overcompensate in their joystick movements, and Day concluded that the effect of latency is typically an oscillation in the operator control response.[3, 4]

These studies also noted that the finer and more complex the human control response required, the worse the effect of latency on operator performance. Teleoperating a weapon mounted on a mobile platform to track a moving target from a potentially moving platform is a significantly more complex task than merely teleoperating a vehicle. This is particularly true for aiming at a distant target. This research suggests that communication latency could have devastating effects on teleoperation performance of a weapon payload when tracking moving targets. A study of the effects of latency on simulated precision shooting in the first-person-shooter game Unreal Tournament 2003 was performed, and concluded that, "…precision shooting is very sensitive to latency, with a decrease in hit accuracy for latencies of 100ms or over."[5]

1.3 Limited Situational Awareness

Most currently developed weapons payloads carry multiple visual sensors. The typical weapon payload carries a wide-angle "scene" camera, a telescopic bore-sight camera, and, possibly, other sensors such as FLIR or radar. This paper will focus on the visual cameras, which tend to be the primary sensors for most UGV weapon payloads.

The scene and bore-sight cameras allow a weapon payload operator to mimic the actions of a human sniper. The sniper uses his unaided eyes or a telescope to view a scene for a potential target. Upon target identification, the sniper switches to a more powerful rifle scope to precisely aim the weapon. Similarly, a UGV operator uses the scene camera to scan for targets and initiate the targeting process, then uses the bore-sight camera to obtain a precise aim. However, these two camera feeds have several limitations. First, the image sizes and resolutions displayed at the operator are limited by the throughput of the communication link and the available screen real-estate. Often both cannot be displayed simultaneously without further reducing their size. If they are displayed one at a time, then the user must manually switch between the two modes during the targeting process. Yanco and Drury performed a study of robot operator situational awareness and concluded that the robot operators in their experiment spent a large amount of time exclusively devoted to acquiring situational awareness (30% on average), sometimes ignoring other critical feedback to the detriment of overall performance. [6]

1.4 Suggested Approach

This limited situational awareness, in combination with the problems of operator burden and latency, should be taken into account in the design of a weapon payload and its associated user interface. SSC San Diego has taken the approach that the majority of the burden of target acquisition (aiming) should be assumed by machine intelligence located on the UGV itself. The human operator, however, retains the task of final target verification, and the arming and firing of the weapon. This approach greatly reduces the negative effects listed above.

For example, the burden of the operator is greatly reduced. The user need only to designate targets through a point-and-click interface, the weapon movements in targeting being automated by an onboard processor. This allows the operator to focus on target designation and the rules of engagement, not on the difficult task of targeting.

Because all the targeting and tracking processing takes place on an onboard processor and all the sensors feed directly to this process, all network latency is also removed from the targeting process. In addition, all other systematic latency can be modeled explicitly in the control software so that the control system can respond to latency in an optimal manner, without the oscillations exhibited in the human control response.

An automated system could also improve an operator's situational awareness in two ways. First, relieved of the duties of manually controlling the weapon, the operator can focus more attention on the scene, and in the identification of targets. Second, the automated control system can also operate the cameras, and provide the user with a high resolution snapshot of the intended target prior to arming and firing. These capabilities, in combination, should reduce the probability of operator error in target identification.

These ideas have been implemented and tested in two prototype weapon payloads: the ROBART III weapon system, and the Networked Remotely Operated Weapon System. These prototypes are further described below.

1. DEVELOPMENT PLATFORMS

2.1 *ROBART III*

ROBART III is a test and evaluation research platform custom built at SSC San Diego. ROBART III hosts numerous sensors for navigation and intruder detection: the SICK scanning laser rangefinder, infrared (IR) sensors, Polaroid ultrasonic transducers, passive infrared (PIR) motion detectors, a gyro-stabilized magnetic compass, and a fiber-optic rate gyro. ROBART III's vision system includes a Visual Stone 360-degree omni-directional camera and a Canon pan-tilt-zoom (PTZ) camera. ROBART III's weapon payload consists of a pneumatic Gatling-style six-barrel weapon mounted as the "right arm." The weapon is aimed via a pan-tilt assembly.

Figure 2. A model of *ROBART* III's weapon payload (left) shown alongside *ROBART III. ROBART's* omnidirectional and Canon pan-tilt-zoom cameras can be seen mounted atop the head assembly.

The weapon payload uses the omnidirectional camera as the scene camera, and the pan-tilt-zoom camera as the high resolution targeting camera. Processing is carried out on an embedded Pentium III-class processor located in the head assembly. This processor is capable of digitizing and analyzing multiple simultaneous video streams in real time and also has access to all of ROBART III's sensor and actuator functionalities.

2.2 Networked Remotely Operated Weapon System

The Networked Remotely Operated Weapon System (NROWS) is a weapon platform specifically designed for automated use, either as a standalone system or mounted on a UGV. NROWS is based on the Telerobotics Corporation (TRC) remote weapon platform, which provides a high-speed, precise aiming platform capable of accepting a wide range of standard light arms. This weapon platform is much faster than all known pan-tilt platforms currently being used as UGV payloads, capable of moving at a rate of 90 degrees per second. The SSC San Diego NROWS prototype is currently configured to carry a standard M4 carbine. A replica M4 Airsoft rifle is used in place of the actual weapon in most testing and development. A scene camera and bore-sight camera provide visual feedback to the user and vision processing system.

Figure 3. The NROWS System mounted on the MDARS UGV (right), and depicted in a wall-mount configuration.

NROWS can be controlled by either wired or wireless networks. All communication to and from NROWS, including multiple live video streams, is carried over standard Ethernet networking, though other networking solutions are possible. NROWS is also designed for easy coordination of multiple weapons. A single operator can assume control and monitor the status of multiple simultaneous weapons. This capability is particularly attractive in a physical security application where multiple weapon payloads are guarding a fixed asset.

As in the *ROBART III* weapon payload, visual sensor data is processed by a Pentium-III class embedded processor that is co-located with the weapon itself. This location greatly reduces the latency involved in transmitting live video over a digital network.

2. WEAPON PAYLOAD AUTOMATION

The process of operating a weapon payload has been broken down into three distinct steps: target detection, target acquisition, and target prosecution. These steps are analogous to the steps a human would use in operating a handheld weapon. Varying degrees of weapon payload autonomy can be achieved by automating individual steps, while leaving others steps under manual control. In the experiments presented in this paper, the first step is entirely automated, the second is semi-autonomous (the human decision-making process is aided by computer-generated alerts and additional computer-generated visual information), and the third step is purely teleoperated. Figure 4 outlines the differences in control responsibility between a purely teleoperated weapon payload and that used in the SSC San Diego prototype semi-autonomous weapon payload prototypes

	Teleoperated Payload		**Semi-autonomous Payload**	
Control Responsibility	Automated	Human	Automated	Human
Target Detection		X	X	

Target Recognition		X		X
Target Acquisition		X	X	
Target Tracking		X	X	
Arming		X		X
Firing		X		X

Figure 4. In the proposed semi-autonomous payload, the human retains the responsibility of positive target identification, arming, and firing, while the rest of the process is automated

3.1 Target Detection

Most weapon payload systems use a "scene" camera to provide situational awareness and scene understanding to the operator. This camera typically has a large field of regard to minimize the "blind area" that the operator can't see. The view from this scene camera is used to identify potential targets. The methods of identification can vary, but the primary visual cues used by human operators are motion-based cues and appearance-based cues. A potential target's speed and motion characteristics can often easily distinguish the target from background clutter in the scene camera. Appearance cues are also used for target detection. These cues include such things as the color and texture of the potential target.

While computer vision techniques have not caught up to the human visual system for target detection capability, an automated system has two advantages over a human operator in target detection. First, an automated system can provide constant, full attention to all areas of a scene. Human, in contrast, cannot sustain attention for long periods of time without suffering from fatigue and reduced performance levels. Second, an automated system can be designed so that it sustains the same level of detection performance regardless of the number of targets that appear in the scene. For example, if 20 targets were to simultaneously appear in the scene, a human might have trouble quickly locating and counting all of them, while an automated system would have no problem with that task.

Furthermore, a human and automated system can "team" together to complement each other's strengths. For example, an automated system could alert a human operator to the presence of potential targets and overlay their locations over live video. Then, the human operator could use his superior detection skills to verify or dismiss the target.

In the implementations described here, all detected targets are tracked continuously while they remain in the view of the scene camera. However, no action is performed. The action of recognizing a target, making a decision and aiming the weapon is left to the target acquisition step in the weapon payload operation process.

It should be noted that the term "detection," in this paper, means the detection of possible targets within a scene, as opposed to target "recognition," which usually means the positive identification of an individual target. In the case of "detection" the target has not been classified or positively identified.

3.2 Target Acquisition

Target acquisition is the process of selecting a target from a list of targets detected in the previous stage and then aiming the weapon at the target. This includes maintaining the aim as the target or UGV continues to move after the gun is brought to bear on the target. There are two primary technical challenges in target acquisition: coordinate system calibration and smooth accurate tracking in the presence of possibly noisy visual tracking data.

First, the target detection step above typically involves tracking objects in the 2D image coordinates of the scene camera. However, this 2D information is not enough to accurately aim a weapon payload at a target. Typically a calibration of some sort is needed to convert image coordinates to world coordinates. In addition, an estimate of the distance to the target is needed to achieve accurate target acquisition. The distance to a target cannot be directly calculated from a monocular scene camera. The implementations below describe two different sets of heuristics and assumptions which achieve this camera calibration with sufficient accuracy for many applications, unassisted by radar or ladar sensors.

Second, the results of computer vision algorithms are often noisy. Specifically, the calculated locations of detected targets that are calculated from incoming image streams can often vary significantly from the true location values and are not sufficient enough for the smooth, accurate aiming of a weapon payload at the tracked target. The implementations in section 4 describe methods of overcoming this noise.

While the choice of which target to aim at is trivial in the case where there is only one detected target, the decision can become very complex in some scenarios. In the case of the presence of multiple potential targets, there are many factors which can affect the "correct" decision. These include such things as target proximity to some asset, target proximity to the UGV, likelihood of target escape, or the perceived threat level of the target. While the implementations described below have only begun to explore this decision-making process, the system architecture has been designed for the future inclusion of one or more decision-making algorithms. The current implementation allows complete human control of target designation, or automatic designation based on one of three criteria: 1) time of detection (aim at first detected target first), 2) proximity to weapon (closest target first), and 3) minimum gun movement (quickest "kill" first)

3.3 Target Prosecution

Target prosecution is defined, in this paper, as the arming and firing of the weapon payload. While there are some applications and scenarios where the automation of target prosecution may be appropriate in the implementations described in this paper, these functions remain entirely under manual, teleoperated control. This is primarily for safety purposes.

4. IMPLEMENTATION

4.1 *ROBART III*

ROBART III's automated weapon payload system includes an appearance-based target detection system, with a unique laser guidance system for target acquisition. The payload can autonomously detect and track multiple targets, and then rapidly acquire them in sequence at the direction of a human operator.

4.1.1 Target Detection

ROBART III's target detection system takes conventional digital images of targets as input. The digital images are used to target templates consisting of features calculated from the input image. These features are matched against incoming images from either the 360-degree camera or rectilinear camera mounted on *ROBART III*'s head assembly. A probabilistic map of potential matching targets is created in real-time for each image in the target templates. Any match which exceeds a preset match threshold is designated as a detected target. Current templates include simple objects such as soda cans.

There are two primary matching algorithms which determine the likelihood of a match. The first is a conventional cross-correlation algorithm which correlates, pixel-by-pixel, the target template over each incoming image as a sliding-window. Image hue is used in the correlation, providing two advantages. First, hue is independent of brightness, making the matching process independent of the lighting condition. Second, correlating on a single channel of data reduces the computational complexity of the process as compared to processing red, green, and blue color channels.

However, cross-correlation is perspective dependent. This means that the object being matched in the scene camera must have a similar scale and orientation as the image of the object being used as the template. Therefore, a second matching algorithm is used simultaneously. This algorithm matches the seven Hu moments between the template image and the incoming image stream. The Hu moments are invariant to scale, rotation, and reflection. The combination of the two algorithms is a robust matching system with a very low occurrence of false positives. The addition of a matching system using so-called "SIFT" features is also being investigated.

4.1.2 Target Acquisition

The target acquisition strategy for *ROBART III* is simply to aim at and track the detected target with the strongest match. However, other acquisition strategies can easily be added to the system. The current target acquisition process for aiming ROBART *III*'s arm-mounted weapon at the target is a two stage process. The first stage involves panning the weapon to roughly the direction of the target. The second stage employs a unique laser-targeting system to precisely aim at the target that overcomes the camera calibration problem presented in 3.2.

The first stage uses a rough calibration between one of *ROBART* III's two scene cameras and the weapon's pan-axis. The two scene cameras, the omnidirectional visual sensor and the pan-tilt zoom camera, provide the image coordinates of a target. The omnidirectional camera's central axis is parallel to the weapon's pan axis and at a known, fixed distance from the weapon's axis. This fixed geometry allows image coordinates from targets detected in the omnidirectional image space to easily be converted to pan-axis coordinates in the weapon's pan axis space. The image-space location of targets detected in imagery from the pan-tilt-zoom camera is similarly easily converted into weapon pan-axis coordinates. The pan-axis of the pan-tilt-zoom camera is also parallel to the weapon's pan-axis, and at a known, fixed distance. The pan-tilt-zoom camera uses a simple centering algorithm to center the target in its field-of-view. Once the target is centered in image coordinates, the pan-coordinates are read from the camera's pan-motor encoder and can be used to pan the weapon to roughly the same orientation.

This process is sufficient to put the weapon within ~10 degrees of the correct pan position if the target is within approximately 20 meters of the target. This is close enough to permit the second stage of the aiming process to take over.

During the second stage, a bore-sight laser sighted along the weapon's active barrel is turned off in synchronization with the vision system's frame capture system and at one half the frequency of image capture. This allows simple image differencing to very accurately locate the laser dot in image space. The laser is simple and low powered, similar to those used as conference room pointers, and is not capable of measuring distances.

Once the laser is located, the weapon is panned and tilted small, fixed distances. Then the laser is re-located in the scene camera's image space. This quick calibration stage provides a relationship between the image space and the pan-and-tilt space of the weapon and allows the weapon to very quickly be aimed at any point in the current field-of-view of the image. The calibration can become invalid if the distance to the surface reflecting the laser dot varies widely as the weapon moves. However, the calibration is updated several times per second as the weapon moves, and has proven to work robustly in cluttered indoor environments.

Once the calibration has been achieved and both the target and laser dot are in the field of view of the pan-tilt-zoom or omnidirectional camera, the weapon can be pan and tilted directly over the target. A final solution is achieved when the target location and the detected location of the laser dot coincide exactly or within a preset tolerance. This method also works for moving targets and moving UGVs, though extensive testing has not been performed with significant UGV or target motion.

4.1.3 Target Prosecution

Arming and firing of *ROBART III*'s weapon payload is performed via teleoperation from a remote user interface. The user can also easily verify that the target has been hit. Voice feedback from ROBART III provides the user with real-time feedback of the detection and acquisition process, with phrases such as, "Target acquired," etc.

4.2 Networked Remote Operated Weapon System

NROWS employs a motion-based system for target detection. Motion in the field-of-view of the scene camera is identified and tracked, then used to feed a system which makes acquisition decisions, or allows a human operator to make them. NROWS is capable of simultaneously tracking the motion of a large number of moving targets, similar to a radar system. As with *ROBART III*, arming and firing of the weapon payload remains under teleoperated control.

NROWS can easily be adapted to use the same tracking system as *ROBART III*, but its high-speed pan-tilt platform makes it ideal for testing motion-based tracking.

4.2.1 Target Detection

NROWS uses a target detection system that detects all motion within the scene camera. The motion detection is currently performed using a statistical, adaptive "background subtraction" scheme. Currently, this system requires that the camera be stationary for at least 1-2 seconds before reliable motion detection can occur, so the UGV must stop briefly. However, the use of 3D sensors such as stereo or ladar and further development of accurate motion-detection-on-the-move algorithms may eliminate the need for a stationary camera in the near future.

Motion detected in the field-of-view of the scene camera is filtered for noise, and roughly classified based on size, location, and aspect ratio before being classified as potential target detection. For example, the aspect ratio of a detected "blob" can distinguish a human from a dog. All motion calculated to be a potential target is then displayed to the user.

4.2.2 Target Acquisition

The user NROWS views the detected motion as computer graphics overlaid over a live video display. If the user identifies a target, a point-and-click interface is used to direct the weapon to both aim at and maintain aim at the source of the motion. Due to the latencies described in 1.2, the user may not be viewing true real-time video. However, the Kalman-filter based tracking system models the motion of each target and allows the system to minimize the effects of any systematic latency.

For the weapon to aim accurately, a camera calibration scheme, as described in 3.2, is necessary. The calibration scheme used in NROWS requires that the pose of the camera, and therefore the UGV be known, and also makes the assumption that the moving object is touching the ground plane. If the ground is not planar, then a terrain map is also necessary. These assumptions currently limit the application of the NROWS payload to a known, fixed environment such as in a physical security application. However, these limitations can be eliminated through the use of an accurate 3D sensor such as stereo or laser, increasing the range of applications possible for the system.

4.2.3 Target Prosecution

As with *ROBART III*, arming and firing remain under teleoperated control.

5. TESTING AND RESULTS

ROBART III's weapon payload has undergone extensive testing in a demonstration involving the "prosecution" of multiple soda cans at a distance of 10-20 feet. While this demonstration scenario does not yet approach the requirements of a real-world application, *ROBART III* is intended to be a research platform with proof-of-concept functionality.

NROWS, however, has been tested in several realistic scenarios involving up to five human involved in a simulated "intrusion" of a large room. The system effectively tracks and, at the user's direction, acquires any of the five “intruders.” However, no objective performance metric has yet been calculated.

7. CONCLUSION AND FUTURE WORK

The SSC San Diego weapon payload prototypes effectively demonstrate weapon payload functionality that will be required for “real world” use of remotely operated weaponry. There are, however many issues to be explored. Comprehensive testing of sensor and tracking performance in a variety of conditions and applications should be performed in order to quantify system performance. Different sensor modalities should be explored, including infrared, radar, ladar, and stereo vision. And, lastly, all this functionality must be presented to the operator in a manner that is simple, and easy to operate.

8. REFERENCES

1. Everett, H.R., Pacis, E.B, Kogut, G., Farrington, N., and S. Khurana, “Towards a Warfighter’s Associate: Eliminating the Operator Control Unit,” SPIE Proc. 5609: Mobile Robots XVII, Philadelphia, PA, October 26 – 28, 2004.
2. Elliott, E. D., and Eagleson, R. (October 1997). Web-based teleoperated Systems using EAI, *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics, Oct 12-15 1997* (Vol. 1, pp. 749-754).
3. Kay, Jennifer, and Thorpe, Charles, “STRIPE: Supervised Telerobotics Using Incremental Polygonal-Earth Geometry”, *Proceedings of the 3rd International Conference on Intelligent Autonomous Systems*, Pittsburgh, PA, February 1993.
4. Day, P.N. (1999). An investigation into the effects of delays in visual feedback on real-time system users. In Sasse, M.A. and Johnson, C. (Eds.) (1999) *Human-Computer Interaction – INTERACT ’99*, Oxford, IOS Press, 674-675
5. Beigbeder, T., Coughlan, R., Lusher, C., Plunkett, J., Agu E., and Claypool, M., “The Effects of Loss and Latency on User Performance in Unreal Tournament 2003,” Pro*ceedings of ACM Network and System Support for Games Workshop,* Portland, USA, September 2004.
6. Drury, J. L., Scholtz, J., and Yanco, H. A. (2003). “Awareness in human-robot interactions.” In *Proceedings of the IEEE Conference on Systems, Man and Cybernetics*, Washington, DC, October 2003.